

Abstract of the Invention

The invention provides methods for identifying related data in a directed graph (e.g., an RDF data set). A “first” step—though the steps are not necessarily executed in sequential order—includes identifying (or marking) as related data expressly satisfying the criteria. A “second” step includes identifying as related ancestors of any data identified as related, e.g., in the first step, unless that ancestor conflicts with the criteria. A “third” step of the method is identifying descendents of any data identified, e.g., in the prior steps, unless that descendent conflicts with the criteria or has a certain relationship with the ancestor from which it descends. The methods generates, e.g., as output, an indication of each of the nodes identified as related in the three steps.